Fostering Mixed-Initiative Visual Analytics through AI Guidance

Alex Endert
Associate Professor and Associate Chair
School of Interactive Computing,
Georgia Tech
endert@gatech.edu

Lessons learned from Alex trying to build mixed-initiative VA tools for the last 15+ years…

~~Fostering Mixed-Initiative Visual Analytics through AI Guidance~~

Alex Endert
Associate Professor and Associate Chair
School of Interactive Computing
Georgia Tech
endert@gatech.edu

Georgia Tech

Goals for this talk

- **Reflect** on recent works (from my lab and others) on mixed-initiative VA
- **Share some lessons learned**, punctuated by **successes** and **failures** along the way
- **Foster a discussion**
  - Feel free ask questions, add comments, etc.
  - I'm going to talk high level about topics. Let's chat more throughout the week!

Why **guidance?**

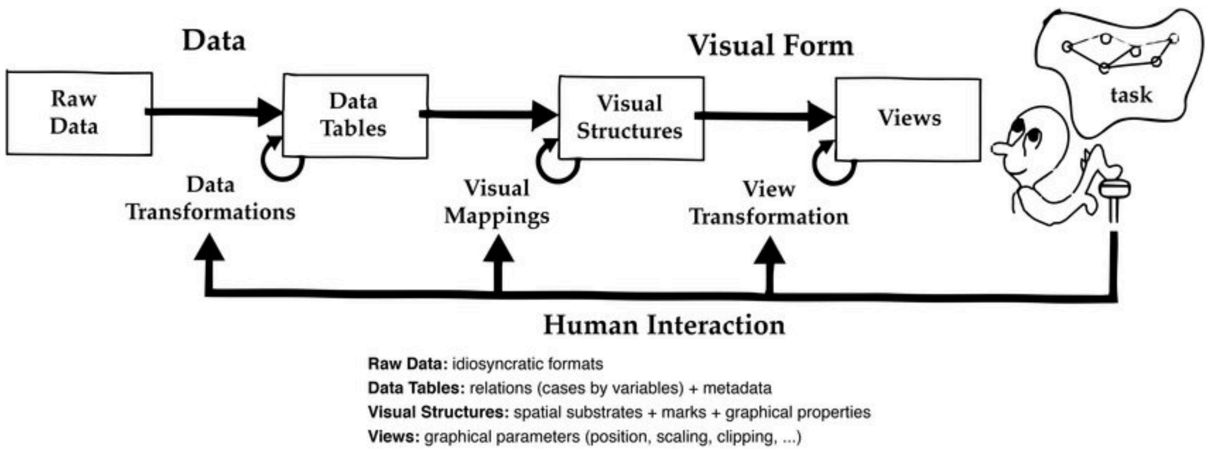

*[Card, Mackinlay, Shneiderman, **1999**]*

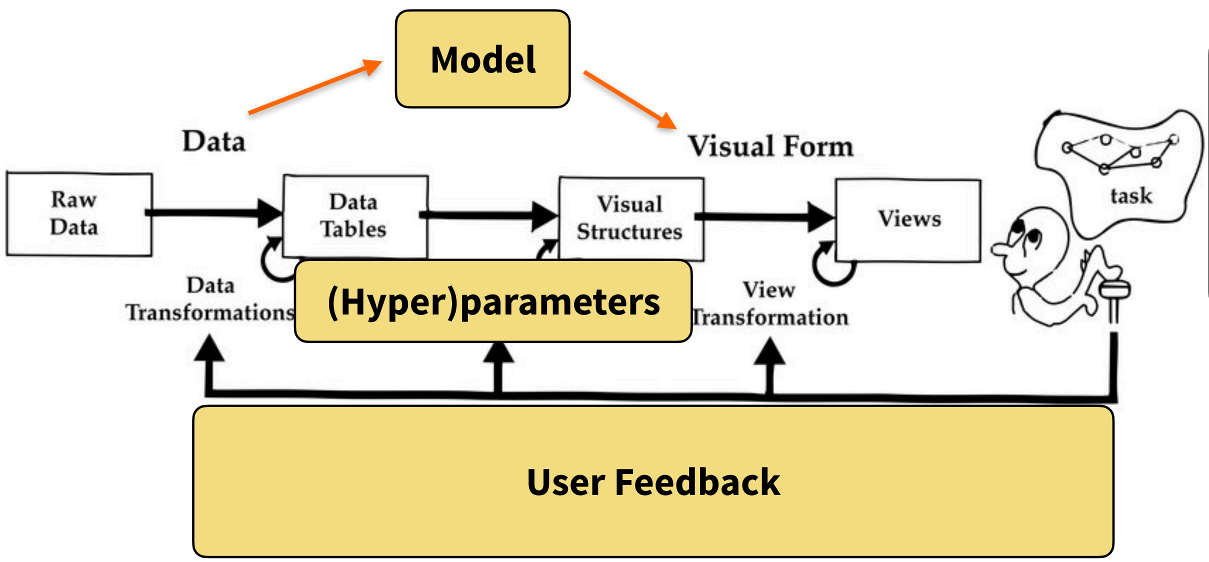

Mental Models
Domain Expertise
Cognitive States/Traits
Cognitive Bias

*[many researchers, ~**2010**]*
*~~[Card, Mackinlay, Shneiderman, **1999**]~~*

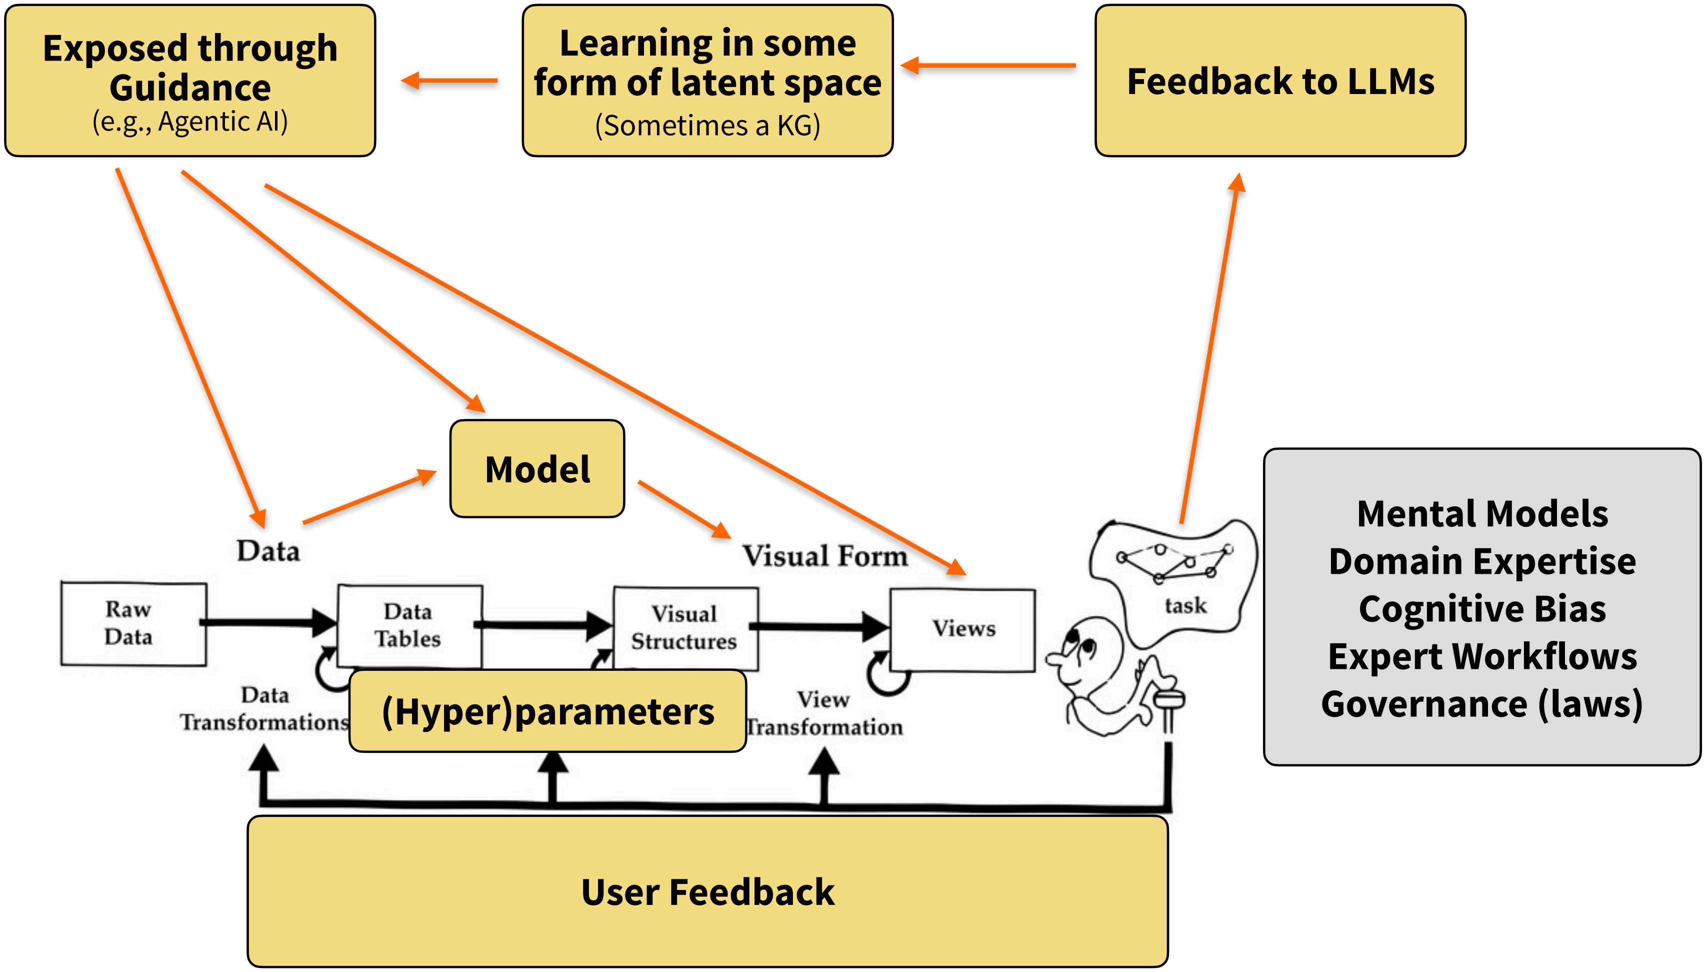

*[even more researchers, ~**2022**]*
~~*[many researchers, ~**2010**]*~~
~~*[Card, Mackinlay, Shneiderman, **1999**]*~~

Why **guidance?**

The opportunity

System are able to learn from people

By analyzing user interactions, systems can **incrementally learn visualization, model, and task specifications**.

Systems can guide/enhance analysis of people

Guidance can help people **accomplish their tasks** better.

LEVELS OF AUTONOMY \ LEVELS OF INTELLIGENCE	RAW DATA	PROCEDURAL	CONTEXT RESPONSIVE	PERSONALIZED	INFERRED INTENT RESPONSIVE	OPERATOR STATE RESPONSIVE	OPERATOR PREDICTIVE
NONE							
INFORMATION FUSER							
SIMPLE AID							
ADVISOR							
INTERACTIVE ADVISOR							
ADAPTIVE ADVISOR							
SERVANT							
ASSISTANT							
ASSOCIATE							
PARTNER							
SUPERVISOR							
AUTONOMOUS							

Figure 3: A taxonomy of two automation levels. For any given system concept, a level of intelligence combined with a level of autonomy describes the system's automation "state".

A GENERAL MODEL OF MIXED-INITIATIVE HUMAN-MACHINE SYSTEMS

Victor Riley
Honeywell Systems and Research Center
Minneapolis, MN

ABSTRACT

The increasing role of automation in human-machine systems requires modelling approaches which are flexible enough to systematically express a large range of automation levels and assist the exploration of a large range of automation issues. A General Model of Mixed-Initiative Human-Machine Systems is described, along with a corresponding automation taxonomy, which: provides a framework for representing human-machine systems over a wide range of complexity; forms the basis of a dynamic, pseudo-mathematical simulation of complex interrelationships between situational and cognitive factors operating in dynamic function allocation decisions; and can guide methodical investigations into the implications of decisions regarding system automation levels.

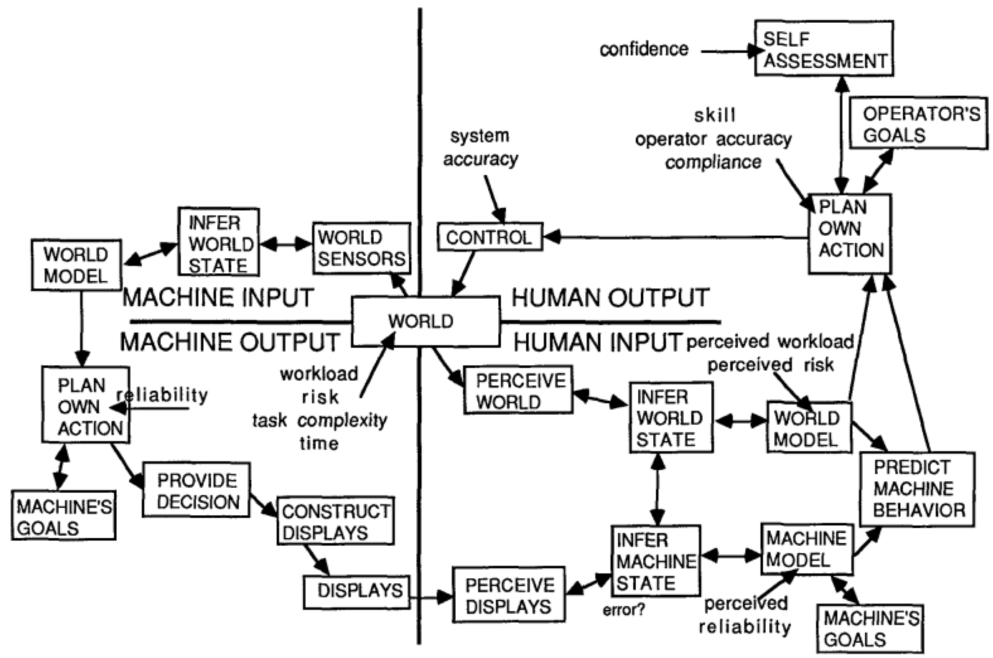

Figure 4. MIM Depiction of a decision aid task. Model parameters are indicated with arrows indicating the components they reside in.

*“Providing mechanisms for efficient agent-user **collaboration**”*

*“Considering **uncertainty** about a user’s goals”*

*“Employing **socially appropriate** behaviors for agent-user interaction … that matches social expectations”*

*“Continuing to **learn** by observing”*

Principles of Mixed-Initiative User Interfaces

Eric Horvitz
Microsoft Research
Redmond, WA 98025 USA
+1 425 936 2127
horvitz@microsoft.com

ABSTRACT

Recent debate has centered on the relative promise of focusing user-interface research on developing new metaphors and tools that enhance users' abilities to directly manipulate objects *versus* directing effort toward developing interface agents that provide automation. In this paper, we review principles that show promise for allowing engineers to enhance human–computer interaction through an elegant coupling of automated services with direct manipulation. Key ideas will be highlighted in terms of the LookOut system for scheduling and meeting management.

Keywords

Intelligent agents, direct manipulation, user modeling, probability, decision theory, UI design

INTRODUCTION

There has been debate among researchers about where great opportunities lay for innovating in the realm of human—computer interaction [10]. One group of researchers has expressed enthusiasm for the development and application of new kinds of automated services, often referred to as interface "agents." The efforts of this group center on building machinery for sensing a user's activity and taking automated actions [4,5,6,8,9]. Other researchers have suggested that effort focused on automation might be better expended on exploring new kinds of metaphors and conventions that enhance a user's ability to *directly manipulate* interfaces to access information and invoke services [1,13]. Innovations on both fronts have been fast paced. However, there has been a tendency for a divergence of interests and methodologies versus focused attempts to leverage innovations in both arenas.

We have pursued principles that provide a foundation for integrating research in direct manipulation with work on interface agents. Our goal is to avoid focusing solely on one tack or the other, but to seek valuable synergies between the two areas of investigation. Surely, we should avoid building complex reasoning machinery to patch fundamentally poor designs and metaphors. Likewise, we wish to avoid limiting designs for human–computer interaction to direct manipulation when significant power and efficiencies can be gained with automated reasoning. There is great opportunity for designing innovative user interfaces, and new human–computer interaction modalities by considering, from the ground up, designs that take advantage of the power of direct manipulation and potentially valuable automated reasoning [2].

PRINCIPLES FOR MIXED-INITIATIVE UI

Key problems with the use of agents in interfaces include poor guessing about the goals and needs of users, inadequate consideration of the costs and benefits of automated action, poor timing of action, and inadequate attention to opportunities that allow a user to guide the invocation of automated services and to refine potentially suboptimal results of automated analyses. In particular, little effort has been expended on designing for a *mixed-initiative* approach to solving a user's problems—where we assume that intelligent services and users may often collaborate efficiently to achieve the user's goals.

Critical factors for the effective integration of automated services with direct manipulation interfaces include:

(1) **Developing significant value-added automation.** It is important to provide automated services that provide *genuine value* over solutions attainable with direct manipulation.

(2) **Considering uncertainty about a user's goals.** Computers are often uncertain about the goals and current the focus of attention of a user. In many cases, systems can benefit by employing machinery for inferring and exploiting the uncertainty about a user's intentions and focus.

(3) **Considering the status of a user's attention in the timing of services.** The nature and timing of automated services and alerts can be a critical factor in the costs and benefits of actions. Agents should employ models of the attention of users and consider the *costs and benefits of deferring action* to a time when action will be less distracting.

(4) **Inferring ideal action in light of costs, benefits, and uncertainties.** Automated actions taken under uncertainty in a user's goals and attention are associated with context-dependent costs and benefits.

[an agent for decision-making]
should be:

Simple
Robust
Easy to control
Adaptive
Complete on important issues
Easy to communicate with

MANAGEMENT SCIENCE
Vol. 16, No. 8, April, 1970
Printed in U.S.A.

MODELS AND MANAGERS: THE CONCEPT OF A DECISION CALCULUS*

JOHN D. C. LITTLE†

Sloan School of Management, M.I.T.

A manager tries to put together the various resources under his control into an activity that achieves his objectives. A model of his operation can assist him but probably will not unless it meets certain requirements. A model that is to be used by a manager should be simple, robust, easy to control, adaptive, as complete as possible, and easy to communicate with. By simple is meant easy to understand; by robust, hard to get absurd answers from; by easy to control, that the user knows what input data would be required to produce desired output answers; adaptive means that the model can be adjusted as new information is acquired; completeness implies that important phenomena will be included even if they require judgmental estimates of their effect; and, finally, easy to communicate with means that the manager can quickly and easily change inputs and obtain and understand the outputs.

Such a model consists of a set of numerical procedures for processing data and judgments to assist managerial decision making and so will be called a decision calculus. An example from marketing is described. It is an on-line model for use by product managers on advertising budgeting questions. The model is currently in trial use by several product managers.

[Little, 1970]

My personal favorite line in the paper when talking about conflicts … :)

[an age
should

Simple
Robust
Easy to
Adaptiv
Comple
Easy to communicate with

(5) *Complete on important issues*. Completeness is in conflict
tures must be found that can handle many phenomena with
important aid to completeness is the incorporation of subject
have a way of making better decisions than their data seem
that they are able to process a variety of inputs and come up wit
about them. So, if you can't lick 'em, join 'em. I say this with
the value of measurement. Many, if not most, of the big advan
edge come from measurement. Nevertheless, at any given po
estimates will be valuable for quantities that are currently
which cannot be measured in the time available before a decis
One problem posed by the use of subjective inputs is tha

OF A DECISION

his control into an
can assist him but
l that is to be used
ive, as complete as
to understand; by
hat the user knows
answers; adaptive
uired; completeness
require judgmental
ith means that the
rstand the outputs.
processing data and
e called a decision
he model for use by
del is currently in

[Little, 1970]

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response

…….. insight!

Interaction → Response
Interaction → Response
→ *Guidance*

Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
Interaction → Response
……… insight!

Interaction → Response

Interaction → Response

→ *Guidance*

…to people to improve limitation of cognitive processes

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

…….. insight!

…to systems to potentially get insight faster

… or perhaps other/additional insight!

Useful to look at through the lens of an

Agentic AI framework

1. **Perceive/Log/Train/Learn**
2. **Reason/Compute**
3. **Act/Guide**
4. **Continue Learning**, iterate

Interaction → Response

Interaction → Response

How to infer user goals and sub-goals?

Guidance

…to people to improve limitation of cognitive processes

How to create and show guidance to users?

How to "skip steps" without "missing insights"?

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

Interaction → Response

……... insight!

…to systems to potentially get insight faster

… or perhaps other/additional insight!

Let's look at a few examples

- **Creation of visuals and charts**
  - Making it easier to create and explore visualizations.

- **Steering and Creating ML Models**
  - Incremental feedback to refine models
  -

Creating Visualizations

predicting and guiding rapid chart creation from natural language interaction and demonstrations

- Translating natural language into into filters, commands, visualizations

NL4DV: A Toolkit for Generating Analytic Specifications for Data Visualization from Natural Language Queries

Arpit Narechania*, Arjun Srinivasan*, and John Stasko

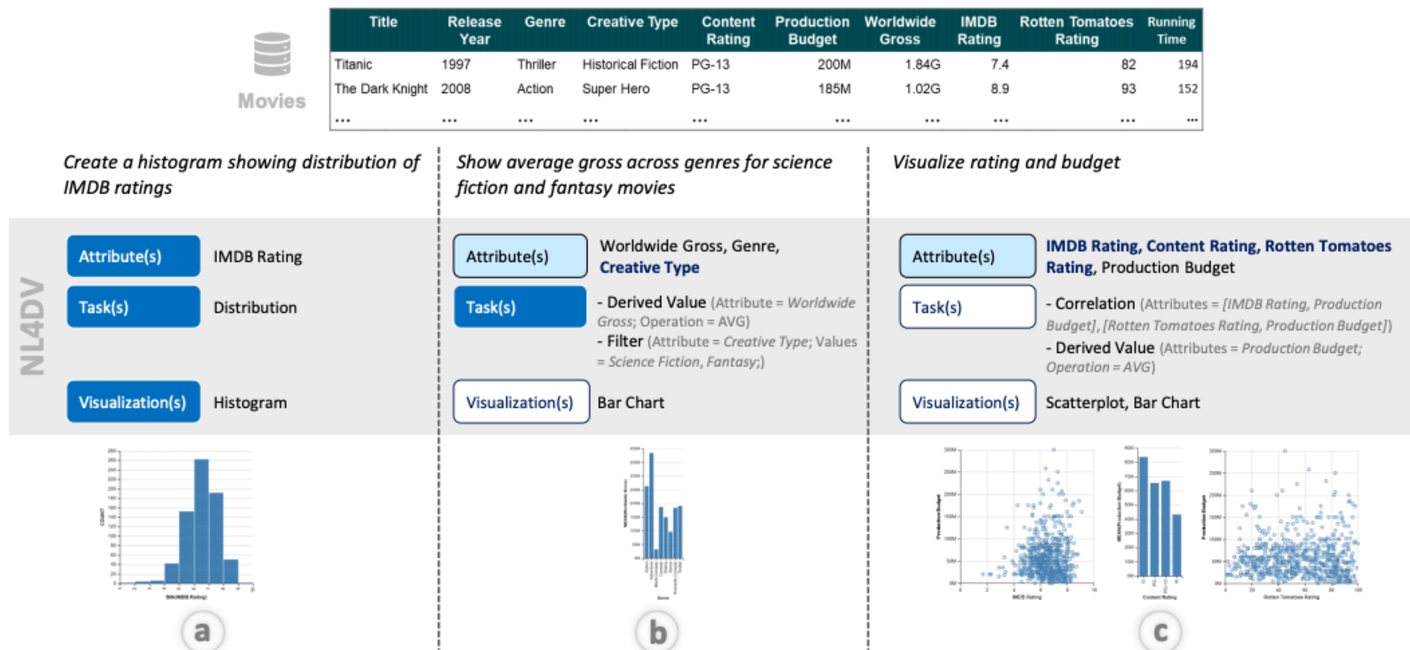

Fig. 1: Examples illustrating the flexibility of natural language queries for specifying data visualizations. NL4DV proces all three query variations, inferring explicit, partially explicit or ambiguous, and implicit references to attributes, tasks, a visualizations. The corresponding visualizations suggested by NL4DV in response to the individual queries are also shown.

https://nl4dv.github.io/nl4dv/

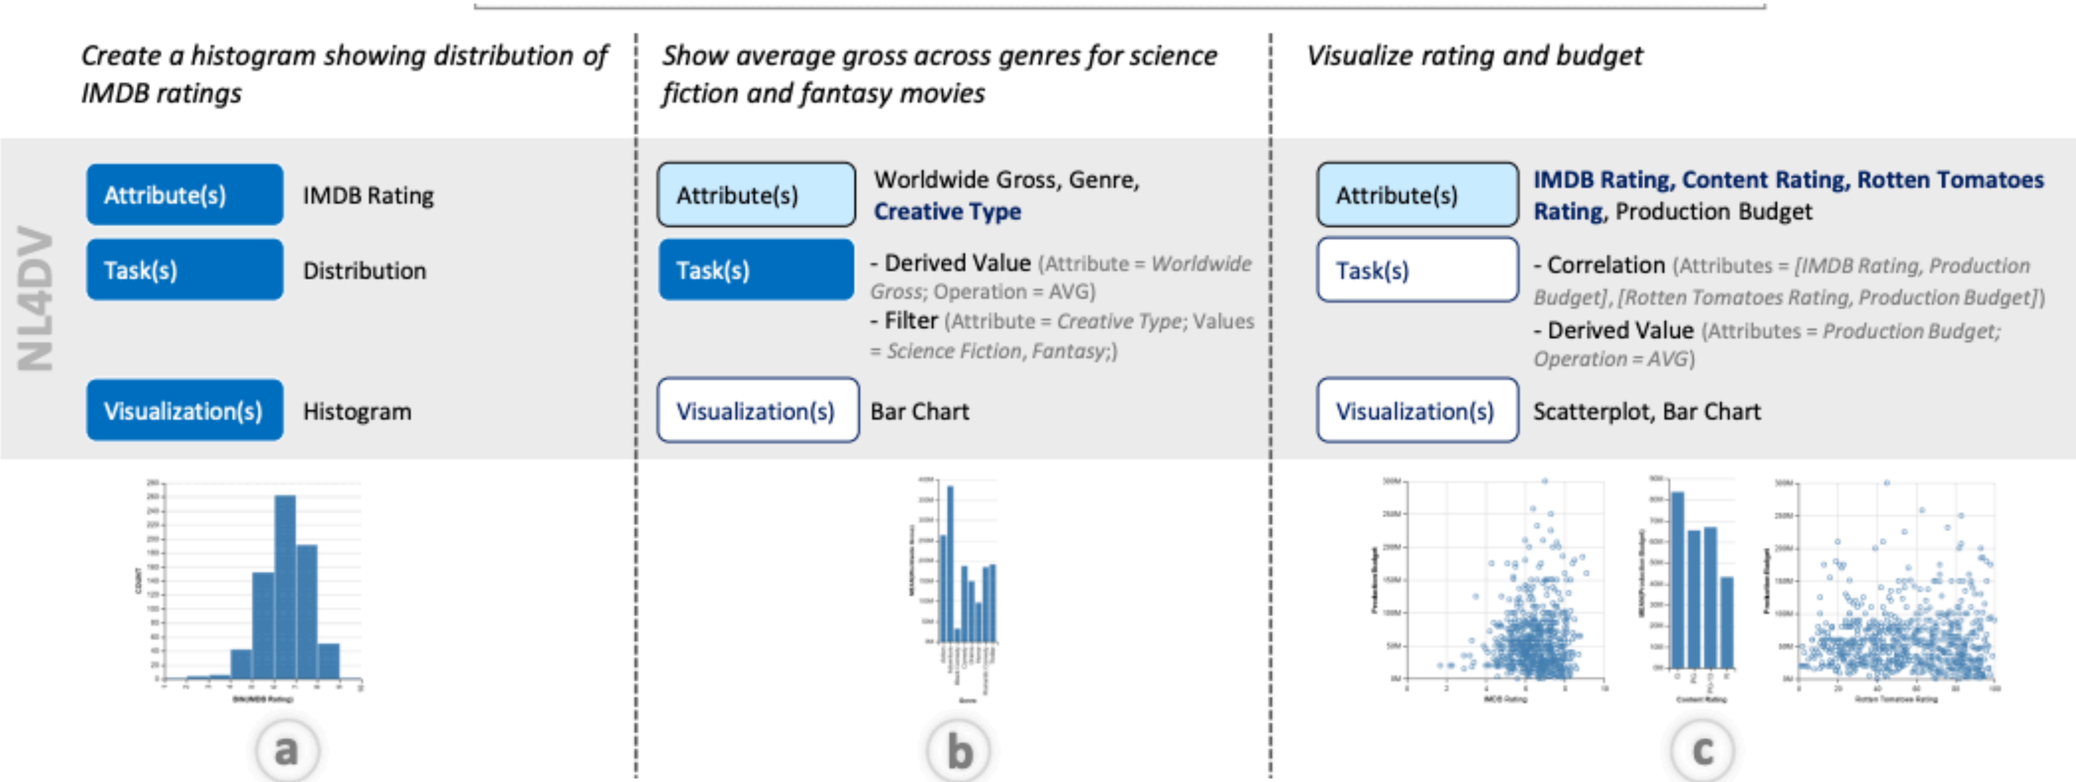
Create a histogram showing distribution of IMDB ratings
Show average gross across genres for science fiction and fantasy movies
Visualize rating and budget
NL4DV
Attribute(s)
IMDB Rating
Task(s)
Distribution
Visualization(s)
Histogram
Attribute(s)
Worldwide Gross, Genre, Creative Type
Task(s)
- Derived Value (Attribute = Worldwide Gross; Operation = AVG)
- Filter (Attribute = Creative Type; Values = Science Fiction, Fantasy;)
Visualization(s)
Bar Chart
Attribute(s)
IMDB Rating, Content Rating, Rotten Tomatoes Rating, Production Budget
Task(s)
- Correlation (Attributes = [IMDB Rating, Production Budget], [Rotten Tomatoes Rating, Production Budget])
- Derived Value (Attributes = Production Budget; Operation = AVG)
Visualization(s)
Scatterplot, Bar Chart
a
b
c

Generating Analytic Specifications for Data Visualization from Natural Language Queries using Large Language Models

Subham Sah[π*]
UNC Charlotte

Rishab Mitra[π†]
Georgia Institute of Technology

Arpit Narechania[π‡]
Georgia Institute of Technology

Alex Endert[§]
Georgia Institute of Technology

John Stasko[¶]
Georgia Institute of Technology

Wenwen Dou[‖]
UNC Charlotte

ABSTRACT

Recently, large language models (LLMs) have shown great promise in translating natural language (NL) queries into visualizations, but their "black-box" nature often limits explainability and debuggability. In response, we present a comprehensive text prompt that, given a tabular dataset and an NL query about the dataset, generates an analytic specification including (detected) data attributes, (inferred) analytic tasks, and (recommended) visualizations. This specification captures key aspects of the query translation process, affording both explainability and debuggability. For instance, it provides mappings from the detected entities to the corresponding phrases in the input query, as well as the specific visual design principles that determined the visualization recommendations. Moreover, unlike prior LLM-based approaches, our prompt supports conversational interaction and ambiguity detection capabilities. In this paper, we detail the iterative process of curating our prompt, present a preliminary performance evaluation using GPT-4, and discuss the strengths and limitations of LLMs at various stages of query translation. The prompt is open-source and integrated into NL4DV, a popular Python-based natural language toolkit for visualization, which can be accessed at **https://nl4dv.github.io**.

Index Terms: Large language models; Natural language inter-

and dataset. However, approaches like NL4DV require developers to create complex rules, which can limit the range and flexibility of input NL queries. Advancements in natural language processing (NLP) and deep learning have further improved NL2VIS systems, which utilize transformers to interpret queries [21, 20].

More recently, large language models (LLMs) like GPT-4 [33], Claude [3], and Gemini [12] have been shown to effectively analyze and extract meaningful information, key concepts, relationships, and trends from unstructured textual data [25]. These capabilities have since been utilized for creative writing [13], code generation [5, 15], dataset curation [18], and visualization creation [6, 34, 9, 41]. One notable LLM-based visualization system, chartGPT [41], has outperformed a parsing-based system (NL4DV [29]) and a deep-learning based system (ncNet [21]). In spite of their superior performance, LLM-based systems have certain documented limitations, such as providing insufficient explanations for the system's generated output [8] and being inconsistent in generating visualizations [23]. These unexplainable, uncertain systems impact transparency and trust, making it difficult for users to find and fix errors. In the NL to SQL domain, several explainable systems have already helped users identify and fix errors in the generated SQL queries [30, 10, 27], motivating this work for more explainable NL2VIS scenarios.

Can we use LLMs to interpret user intent from natural language into analytic specifications?

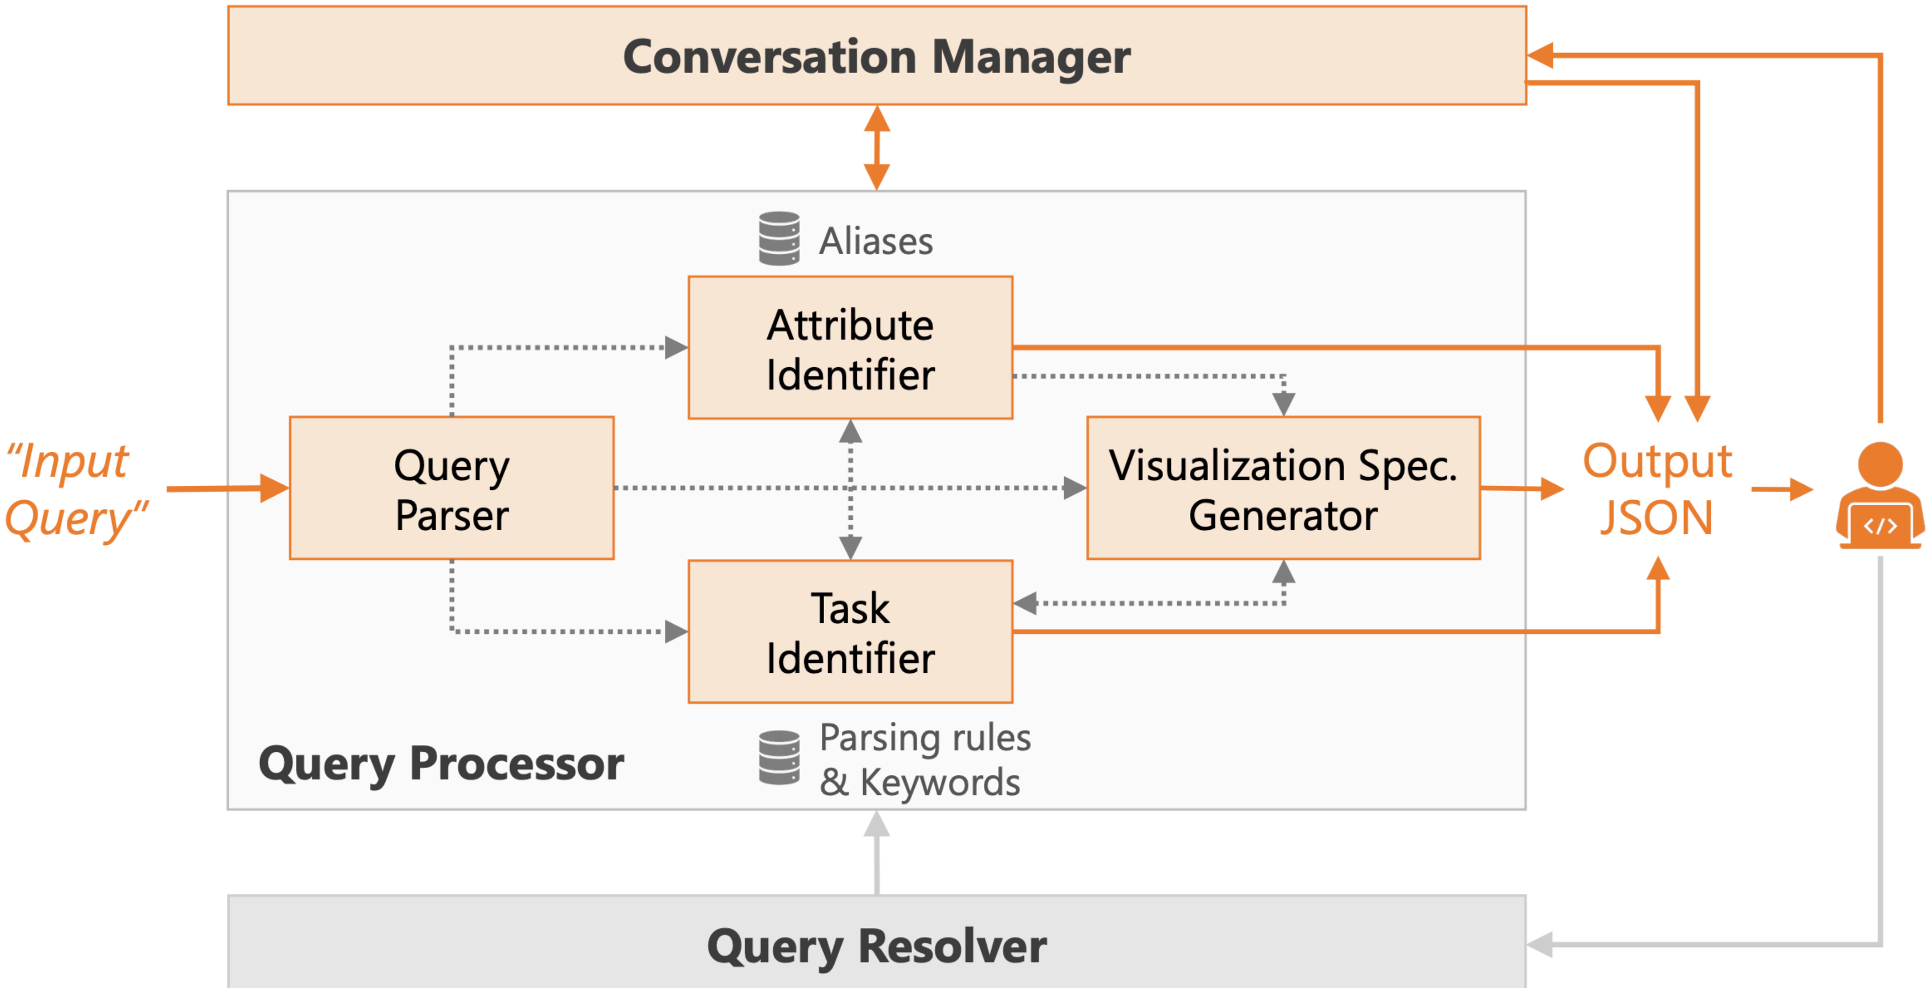
Conversation Manager
Aliases
Attribute
Identifier
"Input
Query"
Query
Parser
Visualization Spec.
Generator
Output
JSON
Task
Identifier
Query Processor
Parsing rules
& Keywords
Query Resolver

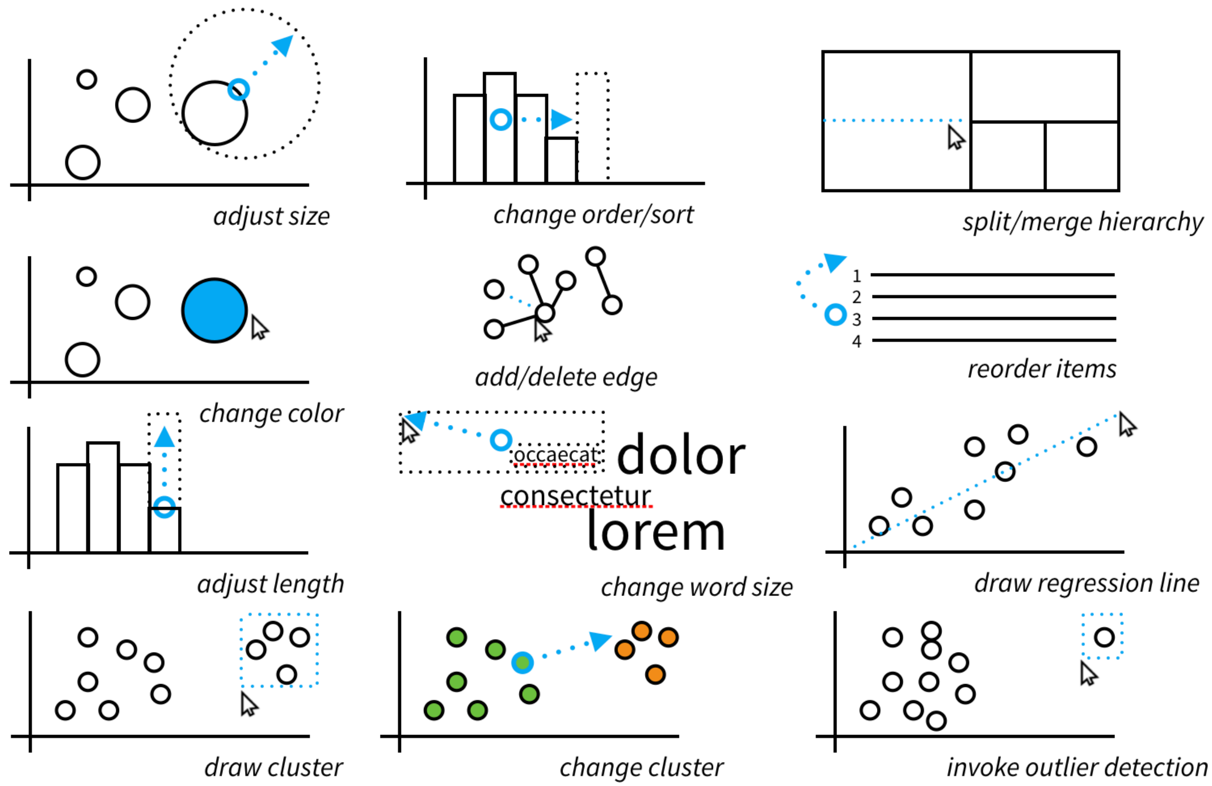

Visualization by Demonstration

Demonstrations to help create visualizations

Saket, B., Kim, H., Brown, E. T., & Endert, A. (2016). Visualization by demonstration: An interaction paradigm for visual data exploration. *IEEE transactions on visualization and computer graphics*, *23*(1), 331-340.

Balance and tradeoff between manual specification and guidance/demonstration

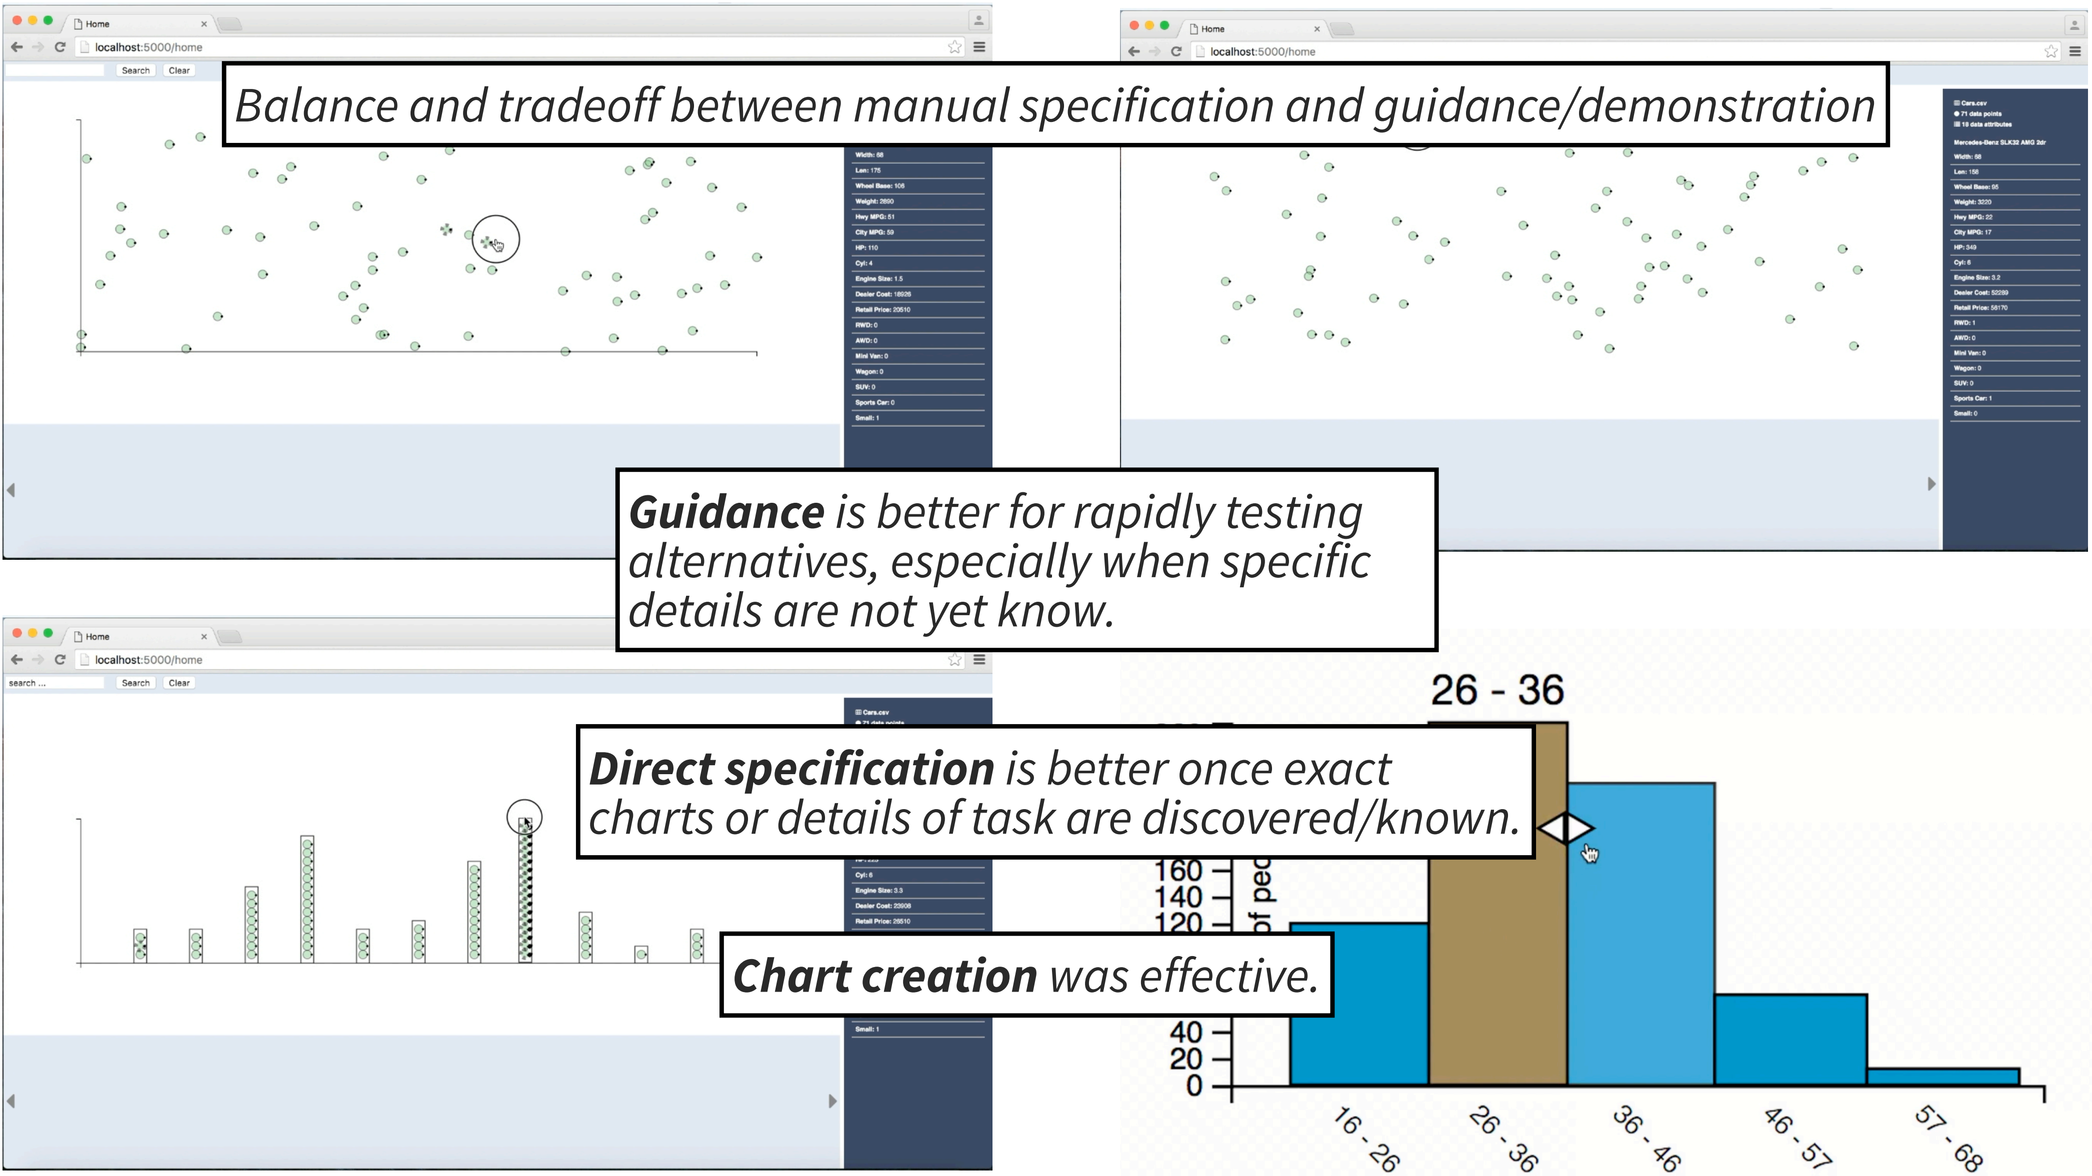

Guidance *is better for rapidly testing alternatives, especially when specific details are not yet know.*

Direct specification *is better once exact charts or details of task are discovered/known.*

Chart creation *was effective.*

Creating and Steering ML and VA Models

Feature and model selection
Graph Modeling
Dimension reduction

Learning from user interaction to guide systems

Model Selection

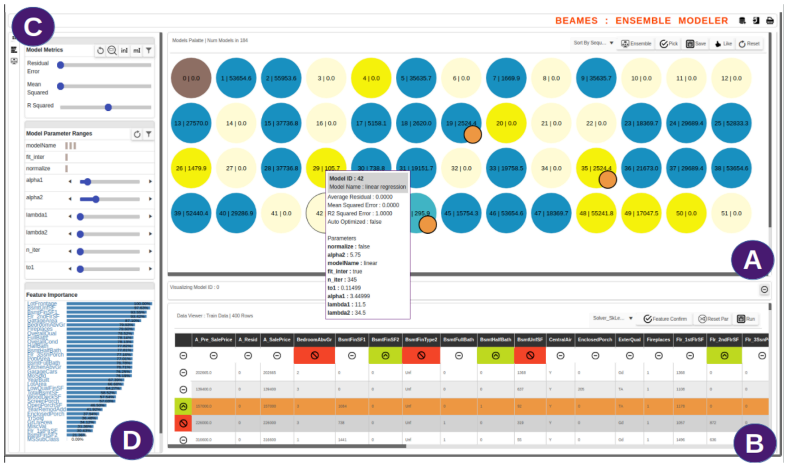

BEAMES: Helping users create, sample, and select regression models


QUESTO: Interactive Construction of Objective Functions for Classification Tasks

Subhajit Das[1], Shenyu Xu,[1] Michael Gleicher,[2] Remco Chang,[3] and Alex Endert[1]

[1]Georgia Institute of Technology, USA [2]University of Wisconsin – Madison, USA [3]Tufts University, USA

QUESTO: Users Express Constraints, System Learns Objective functions for Classification Tasks

Learning from user interaction to guide systems

Model Steering

InterAxis, AxiSketcher, Semantic Interaction: Steering Dimension Reduction models

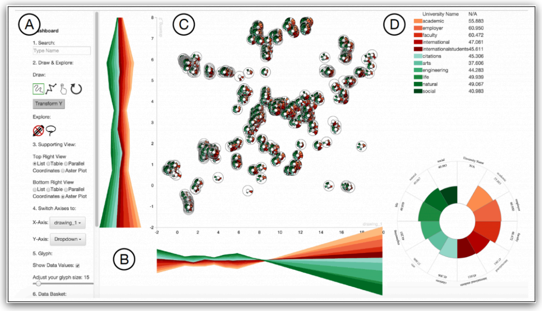

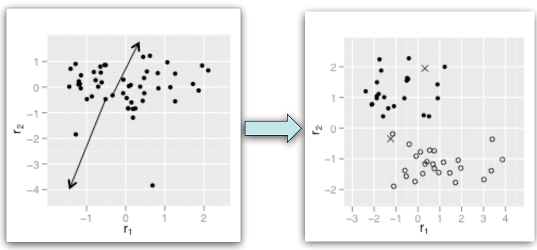

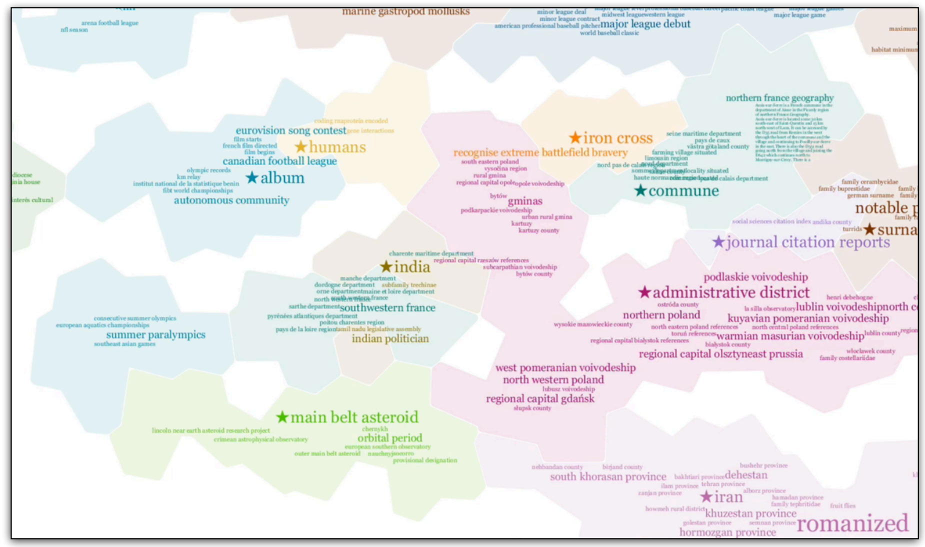

TexTonic: Steering Hierarchical Topic Models

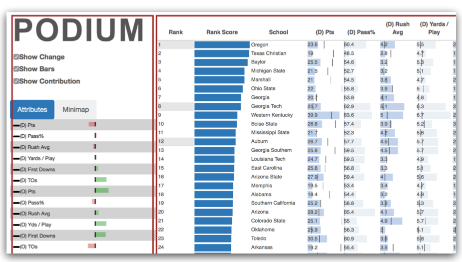

Podium: Steering Tabular Ranking Models

Learning from user interaction to guide systems

Feature and Data Selection

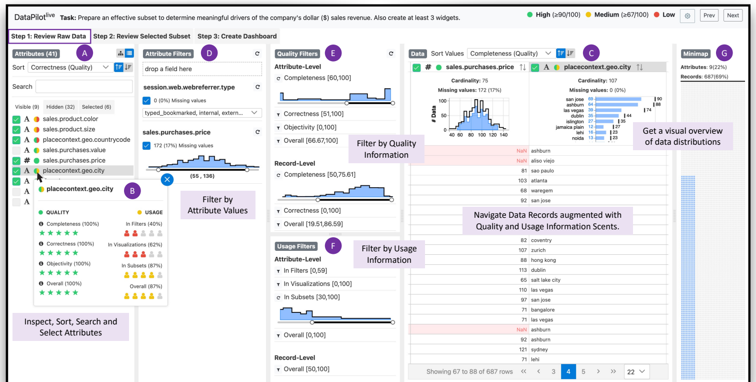

DataPilot: Guidance for data subset and feature selection

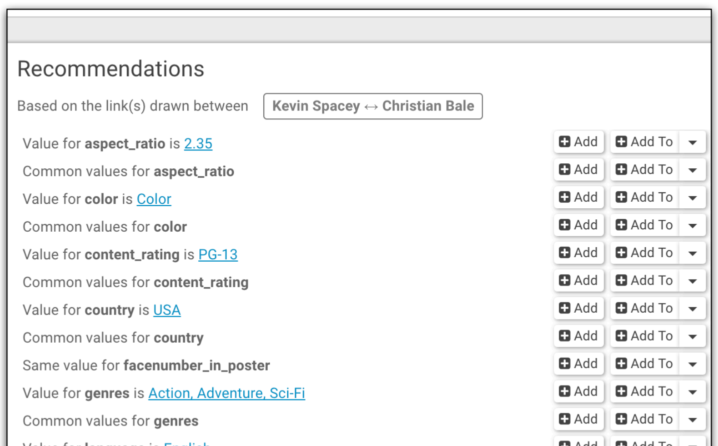

Graphiti: From Tabular Data to Graph Models

What (I think) we got right

Reflection

These tools help users get to a "state" quickly.

They can create visualizations, models. Fantastic.

Replaced direct manipulation of parameters with flexible interactions.

(Mostly) captured user intent and task to produce a result.

Frameworks and libraries created are extensible.

Found generalizable principles across multiple tools.

What (I think) we got wrong

Reflection

Several lessons learned emerged. Let’s go through them.

Lesson 1

The workflow, or process may be more important than the outcome.

Guiding analytic workflow

Using AI to detect and mitigate exploration bias

Lumos: Increasing Awareness of Analytic Behavior during Visual Data Analysis

Arpit Narechania, Adam Coscia, Emily Wall, Alex Endert

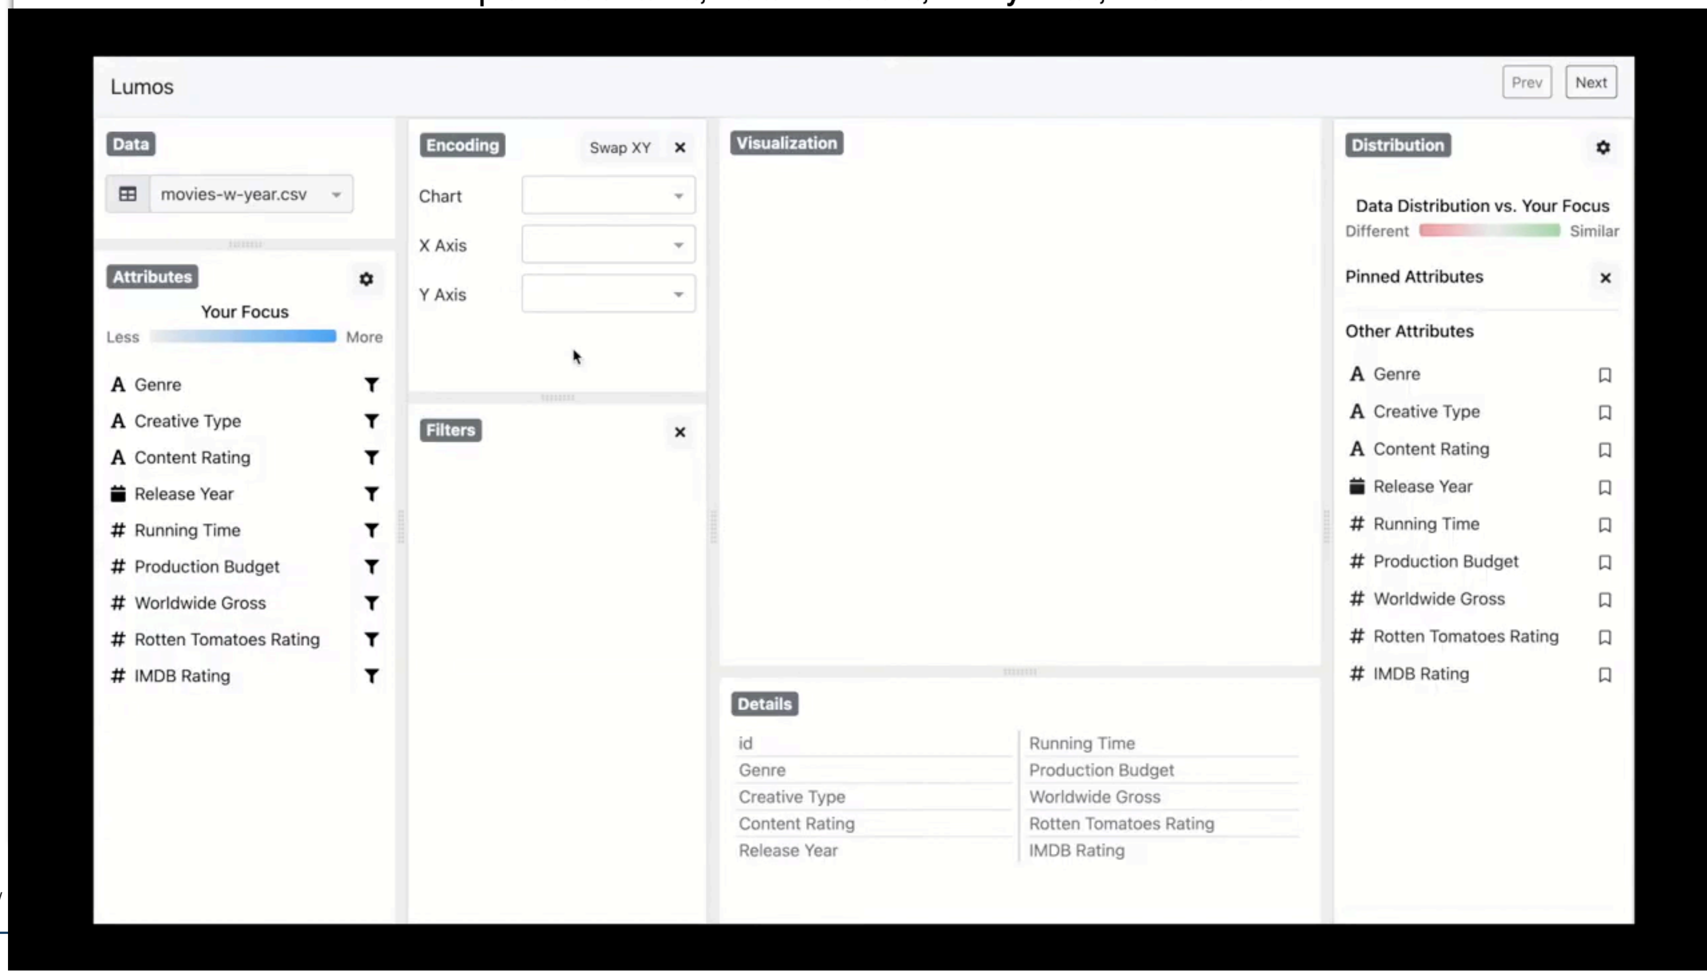

https://lumos-webapp.herokuapp.com/

So, studying this problem in more detail. Is it

showing users their provenance,

or

showing them the system's interpretation of their provenance that's key?

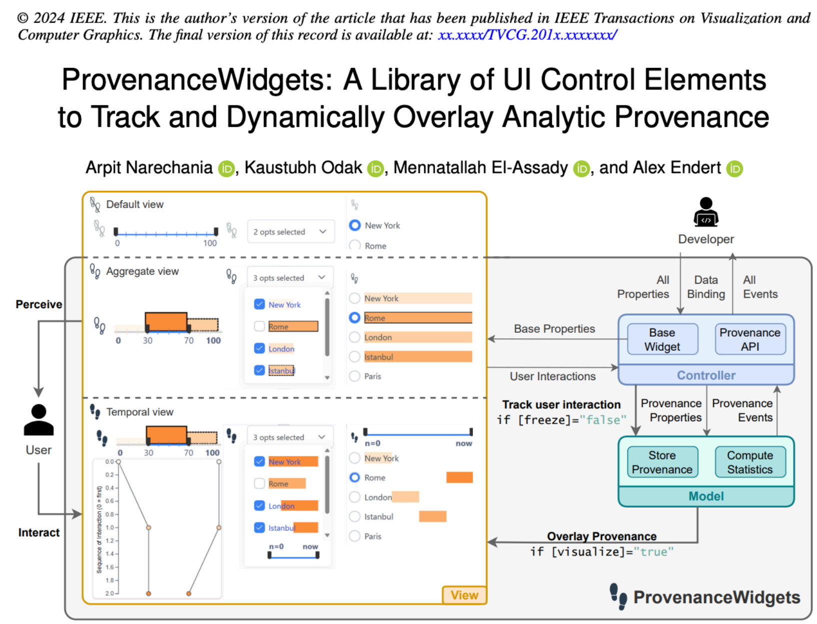

ProvenanceWidgets: A Library of UI Control Elements to Track and Dynamically Overlay Analytic Provenance

Arpit Narechania, Kaustubh Odak, Mennatallah El-Assady, and Alex Endert

Enables developers to directly integrate provenance into standard widgets

We can also use LLMs to make guidance more actionable

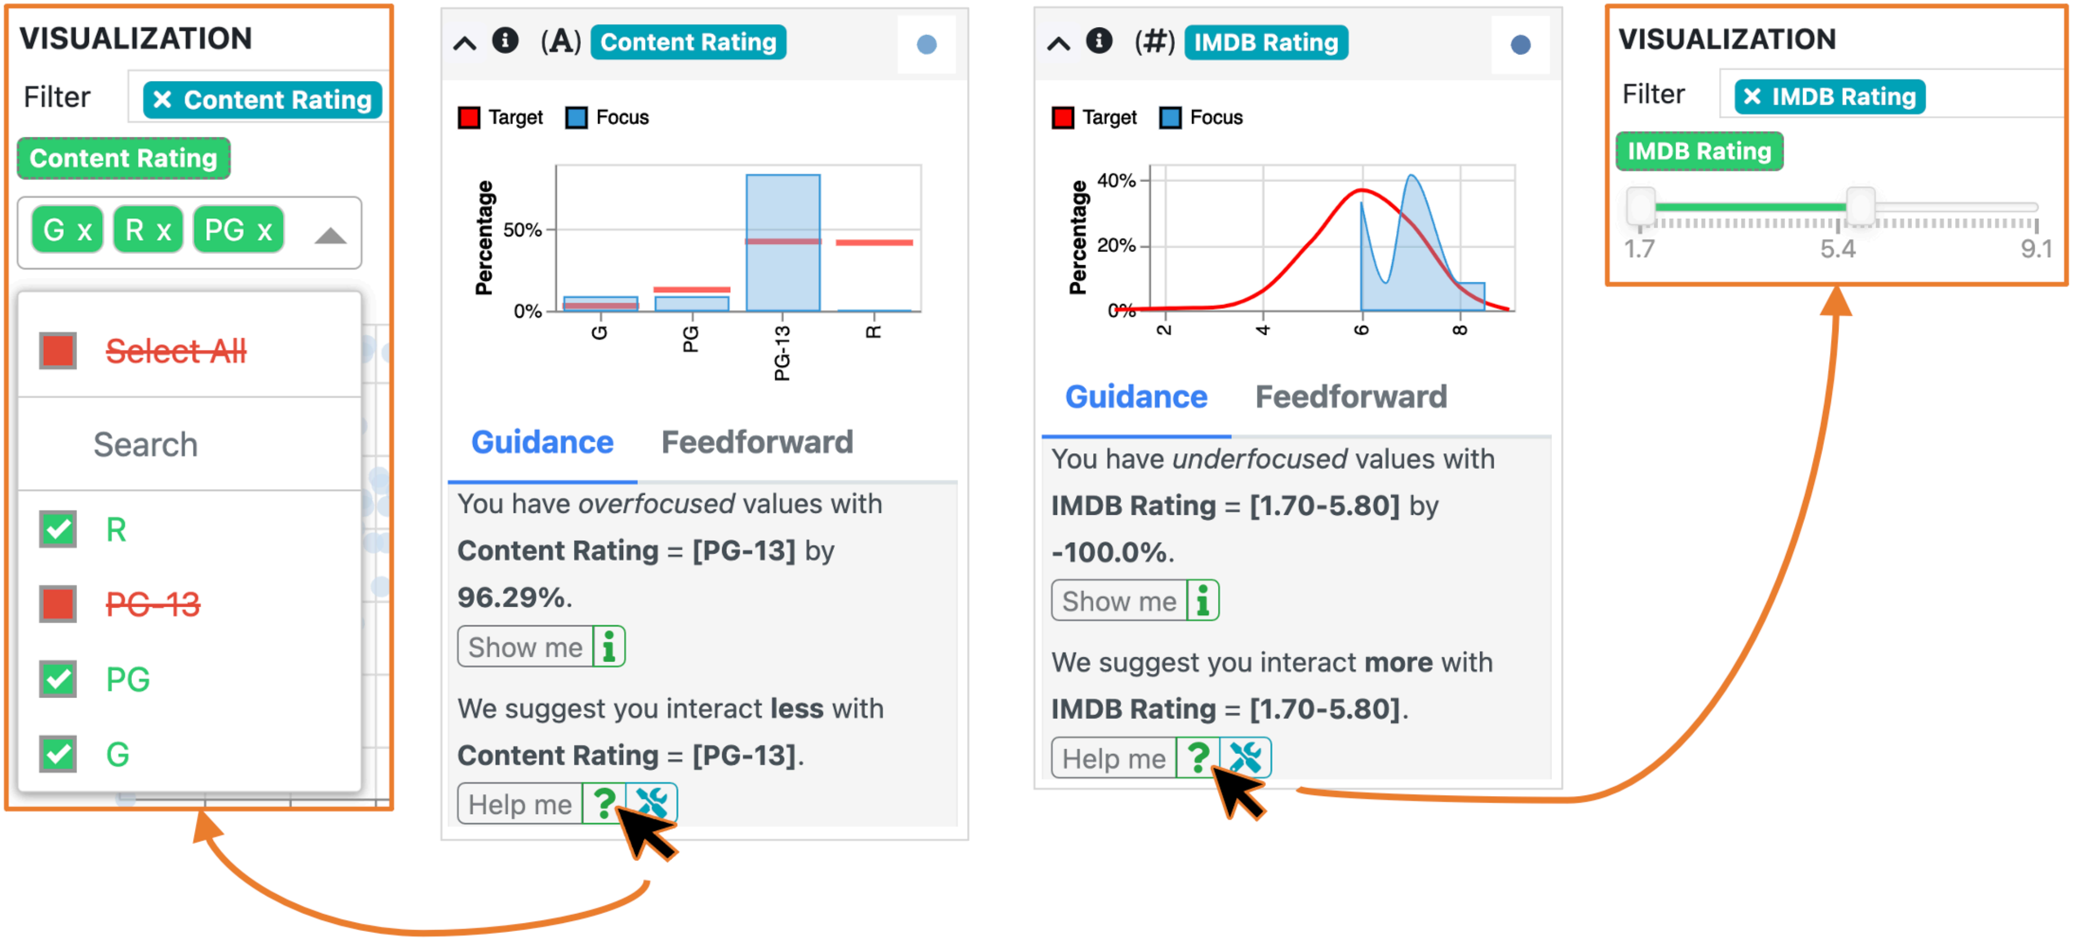

[hot off the press] *Also developing ways to incorporate LLM-generated explainability to guidance*

Lesson 1

The workflow, or process may be more important than the outcome.

Lesson 2

Over-reliance or over-trust may be problematic.

Guidance Source and Over-reliance

- Data subset selection task
- Guidance quality controlled
- Users manually request guidance

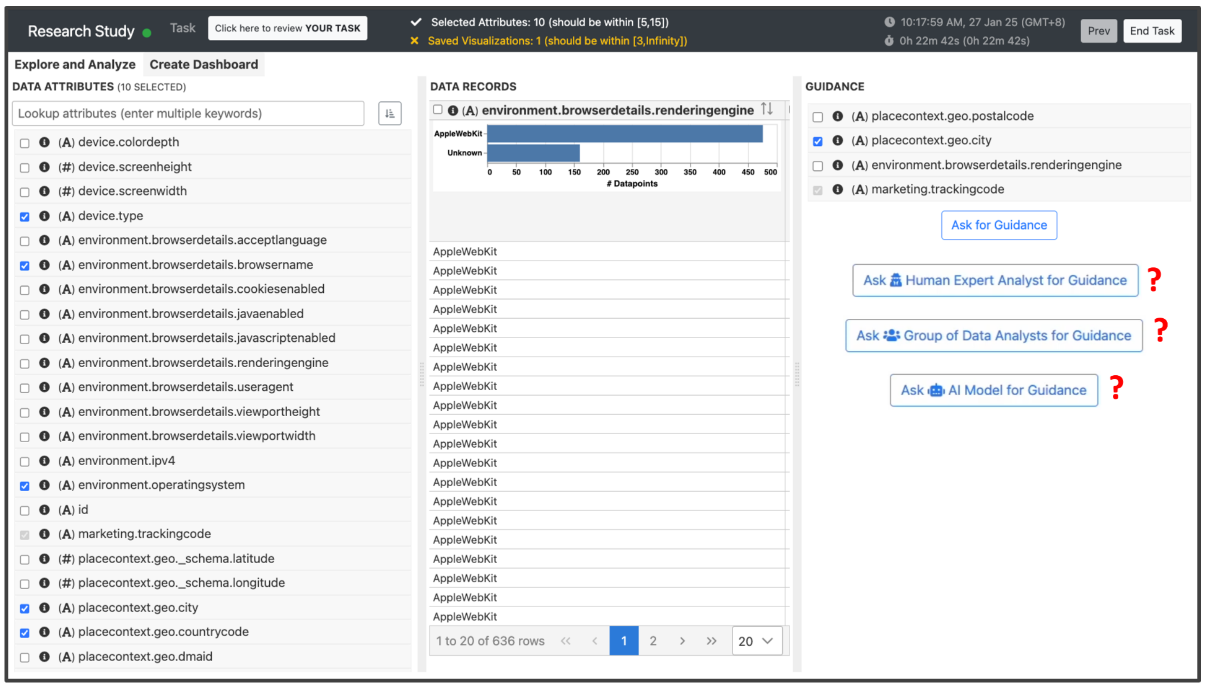

- Findings
  - Users cautiously ask for guidance, but after initial usage begin **blindly trusting it regardless of quality**
  - Dataset understanding suffers

Narechania, A., Endert, A., & Sinha, A. R. (2025, March). Guidance Source Matters: How Guidance from AI, Expert, or a Group of Analysts Impacts Visual Data Preparation and Analysis. In Proceedings of the 30th International Conference on Intelligent User Interfaces (pp. 789-809).

Lesson 1

The workflow, or process may be more important than the outcome.

Lesson 2

Over-reliance or over-trust may be problematic.

Lesson 3

Unit of agency varies by task.

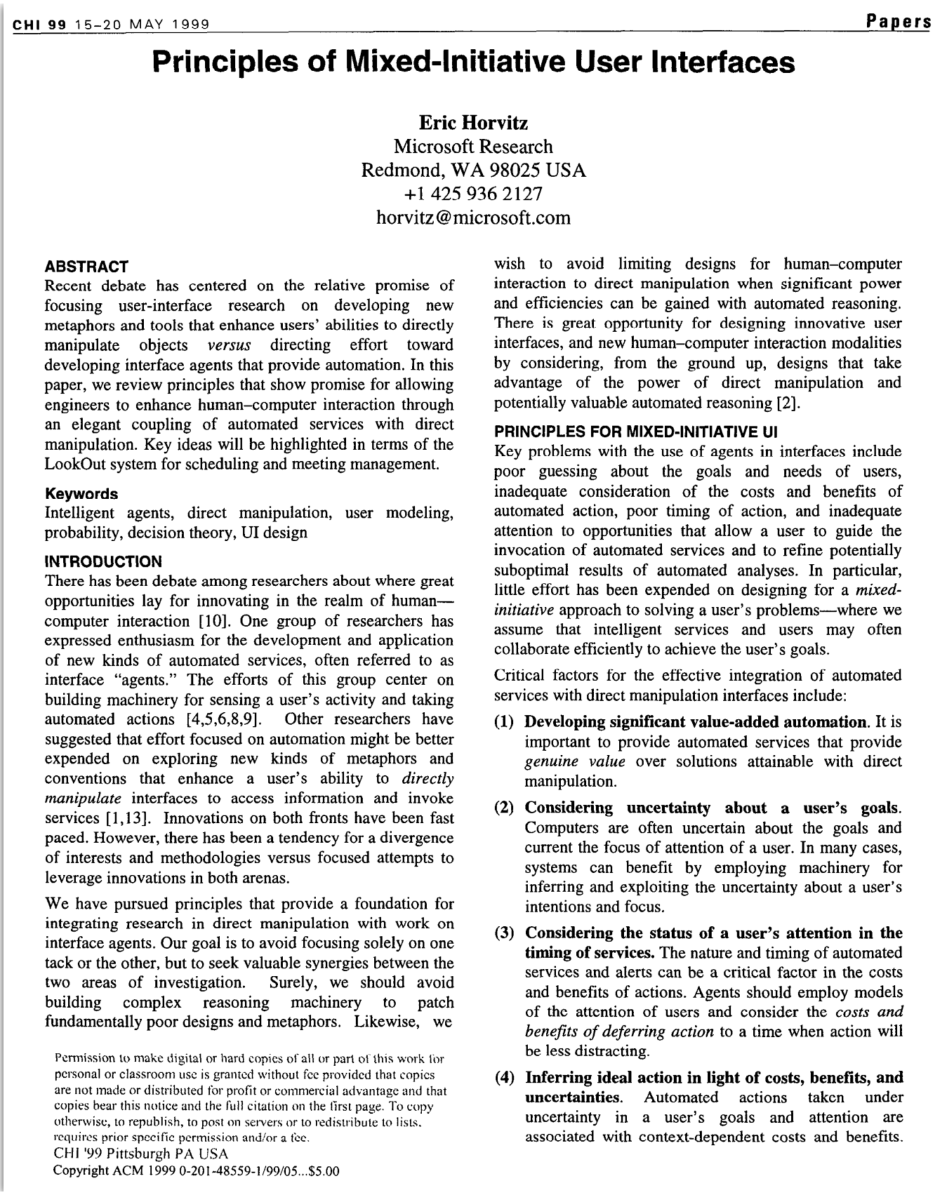

Principles of Mixed-Initiative User Interfaces

Eric Horvitz
Microsoft Research
Redmond, WA 98025 USA
+1 425 936 2127
horvitz@microsoft.com

ABSTRACT
Recent debate has centered on the relative promise of focusing user-interface research on developing new metaphors and tools that enhance users' abilities to directly manipulate objects *versus* directing effort toward developing interface agents that provide automation. In this paper, we review principles that show promise for allowing engineers to enhance human–computer interaction through an elegant coupling of automated services with direct manipulation. Key ideas will be highlighted in terms of the LookOut system for scheduling and meeting management.

Keywords
Intelligent agents, direct manipulation, user modeling, probability, decision theory, UI design

INTRODUCTION
There has been debate among researchers about where great opportunities lay for innovating in the realm of human—computer interaction [10]. One group of researchers has expressed enthusiasm for the development and application of new kinds of automated services, often referred to as interface "agents." The efforts of this group center on building machinery for sensing a user's activity and taking automated actions [4,5,6,8,9]. Other researchers have suggested that effort focused on automation might be better expended on exploring new kinds of metaphors and conventions that enhance a user's ability to *directly manipulate* interfaces to access information and invoke services [1,13]. Innovations on both fronts have been fast paced. However, there has been a tendency for a divergence of interests and methodologies versus focused attempts to leverage innovations in both arenas.

We have pursued principles that provide a foundation for integrating research in direct manipulation with work on interface agents. Our goal is to avoid focusing solely on one tack or the other, but to seek valuable synergies between the two areas of investigation. Surely, we should avoid building complex reasoning machinery to patch fundamentally poor designs and metaphors. Likewise, we wish to avoid limiting designs for human–computer interaction to direct manipulation when significant power and efficiencies can be gained with automated reasoning. There is great opportunity for designing innovative user interfaces, and new human–computer interaction modalities by considering, from the ground up, designs that take advantage of the power of direct manipulation and potentially valuable automated reasoning [2].

PRINCIPLES FOR MIXED-INITIATIVE UI
Key problems with the use of agents in interfaces include poor guessing about the goals and needs of users, inadequate consideration of the costs and benefits of automated action, poor timing of action, and inadequate attention to opportunities that allow a user to guide the invocation of automated services and to refine potentially suboptimal results of automated analyses. In particular, little effort has been expended on designing for a *mixed-initiative* approach to solving a user's problems—where we assume that intelligent services and users may often collaborate efficiently to achieve the user's goals.

Critical factors for the effective integration of automated services with direct manipulation interfaces include:

(1) **Developing significant value-added automation.** It is important to provide automated services that provide *genuine value* over solutions attainable with direct manipulation.

(2) **Considering uncertainty about a user's goals.** Computers are often uncertain about the goals and current the focus of attention of a user. In many cases, systems can benefit by employing machinery for inferring and exploiting the uncertainty about a user's intentions and focus.

(3) **Considering the status of a user's attention in the timing of services.** The nature and timing of automated services and alerts can be a critical factor in the costs and benefits of actions. Agents should employ models of the attention of users and consider the *costs and benefits of deferring action* to a time when action will be less distracting.

(4) **Inferring ideal action in light of costs, benefits, and uncertainties.** Automated actions taken under uncertainty in a user's goals and attention are associated with context-dependent costs and benefits.

Balance (battle?) for agency between people and AI continues,
and may indicate that the **sociotechnical challenges in mixed-initiative systems are crucial** to their success

Alex Endert
Associate Professor
Visual Analytics; HCI; Coffee

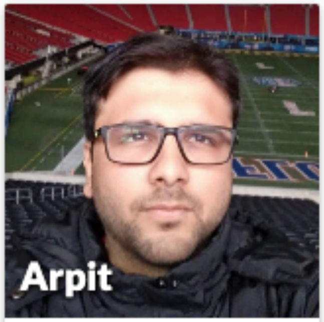

Arpit Narechania
PhD Student, CS
Guidance; Natural Language Processing; Chessmaster

Grace Guo
PhD Student, HCC
Recommenders; Annotations; Musical theatre enthusiast

Adam Coscia
PhD Student, HCC
AI; Language models; Knowledge graphs; Snowboarding

Sichen Jin
PhD Student, CS
Information visualization; Machine Learning

Kaustubh Odak
MS Student, CS
Visual Analytics

Visual Analytics Lab
@ Georgia Tech

SIEMENS
Dolby
DARPA
KEYSIGHT TECHNOLOGIES
Microsoft
GVU GEORGIA TECH
Pacific Northwest NATIONAL LABORATORY
VACCINE
Visual Analytics for Command, Control, and Interoperability Environments
A U.S. Department of Homeland Security Center of Excellence
Ford
NSF
DEPARTMENT OF DEFENSE UNITED STATES OF AMERICA
U.S. DEPARTMENT OF HOMELAND SECURITY
DEPARTMENT OF JUSTICE
charles river analytics
Adobe